arXiv:1406.0893v1 [cs.LO] 3 Jun 2014

Query Answering over Contextualized RDF Knowledge with Forall-Existential Bridge Rules: Attaining Decidability using Acyclicity (full version)

Mathew Joseph[1,2], Gabriel Kuper[2], and Luciano Serafini[1]

[1] DKM, FBK-IRST, Trento, Italy
[2] DISI, University Of Trento, Trento, Italy

{mathew,serafini}@fbk.eu, kuper@disi.unitn.it

Abstract. The recent outburst of context-dependent knowledge on the Semantic Web (SW) has led to the realization of the importance of the quads in the SW community. Quads, which extend a standard RDF triple, by adding a new parameter of the 'context' of an RDF triple, thus informs a reasoner to distinguish between the knowledge in various contexts. Although this distinction separates the triples in an RDF graph into various contexts, and allows the reasoning to be decoupled across various contexts, bridge rules need to be provided for inter-operating the knowledge across these contexts. We call a set of quads together with the bridge rules, a quad-system. In this paper, we discuss the problem of query answering over quad-systems with expressive forall-existential bridge rules. It turns out the query answering over quad-systems is undecidable, in general. We derive a decidable class of quad-systems, namely *context-acyclic* quad-systems, for which query answering can be done using forward chaining. Tight bounds for data and combined complexity of query entailment has been established for the derived class.

Keywords: Contextualized RDF/OWL knowledge, Contextualized Query Answering, Quads, Forall-Existential Rules, Semantic Web, Knowledge Representation.

1 Introduction

One of the major recent changes in the SW community is the transformation from a *triple* to a *quad* as its primary knowledge carrier. As a consequence, more and more triple stores are becoming *quad* stores. Some of the popular quad-stores are 4store[1], Openlink Virtuoso[2], and some of the current popular triple stores like Sesame[3] internally keep track of the context by storing arrays of four names (c, s, p, o) (further denoted as $c : (s, p, o)$), where c is an identifier that stands for the context of the triple (s, p, o). Some of the recent initiatives in this direction have also extended existing formats like N-Triples to N-Quads. The latest Billion triples challenge datasets (BTC 2012) have all been released in the N-Quads format.

[1] http://4store.org
[2] http://virtuoso.openlinksw.com/rdf-quad-store/
[3] http://www.openrdf.org/

One of the main benefits of quads over triples are that they allow users to specify various attributes of meta-knowledge that further qualify knowledge [8], and also allow users to query for this meta knowledge [29]. Examples of these attributes, which are also called *context dimensions* [12], are provenance, creator, intended user, creation time, validity time, geo-location, and topic. Having defined various contexts in which triples are dispersed, one can declare in another meta-context mc, statements such as $mc\colon (c_1, \text{creator, John})$, $mc\colon (c_1, \text{expiryTime, "jun-2013"})$ that talk about the knowledge in context c_1, in this case its creator and expiry time. Another benefit of such a contextualized approach is that it opens possibilities of interesting ways for querying a contextualized knowledge base. For instance, if context c_1 contains knowledge about Football World Cup 2014 and context c_2 about Football Euro Cup 2012. Then the query "who beat Italy in both Euro Cup 2012 and World Cup 2014" can be formalized as the conjunctive query:

$$c_1\colon (x, \text{beat}, \text{Italy}) \wedge c_2\colon (x, \text{beat}, \text{Italy}), \text{where } x \text{ is a variable.}$$

As the knowledge can be separated context wise and simultaneously be fed to separate reasoning engines, this approach increases both efficiency and scalability. Besides the above flexibility, *bridge rules* [4] can be provided for inter-interoperating the knowledge in different contexts. Such rules are primarily of the form:

$$c : \phi(\boldsymbol{x}) \rightarrow c' : \phi'(\boldsymbol{x})$$

where ϕ, ϕ' are both atomic concept (role) symbols, c, c' are contexts. The semantics of such a rule is that if, for any $\boldsymbol{a}$, $\phi(\boldsymbol{a})$ holds in context c, then $\phi'(\boldsymbol{a})$ should hold in context c', where $\boldsymbol{a}$ is a unary/binary vector dependending on whether ϕ, ϕ' are concept/role symbols. Although such bridge rules serve the purpose of specifying knowledge interoperability from a source context c to a target context c', in many practical situations there is the need of interoperating multiple source contexts with multiple target targets, for which the bridge rules of the form (1) is inadequate. Besides, one would also want the ability of creating new values in target contexts for the bridge rules.

In this work, we consider *forall-existential bridge rules* that allows conjunctions and existential quantifiers in them, and hence is more expressive than those, in DDL [4] and McCarthy et al. [27]. A set of quads together with such bridge rules is called a *quad-system*. The main contributions of this work can be summarized as:

1. We provide a basic semantics for contextual reasoning over quad-systems, and study contextualized conjunctive query answering over them. For query answering, we use the notion of a *distributed chase*, which is an extension of a standard *chase* [22, 1] that is widely used in databases and KR for the same.
2. We show that conjunctive query answering over quad-systems, in general, is undecidable. We derive a class of quad-systems called *context acyclic* quad-systems, for which query answering is decidable and can be done by forward chaining. We give both data and combined complexity of conjunctive query entailment for the same.

The paper is structured as follows: In section 2, we formalize the idea of contextualized quad-systems, giving various definitions and notations for setting the background.

In section 3, we formalize the query answering on quad-systems, define notions such as distributed chase that is further used for query answering, and give the undecidability results of query entailment for unrestricted quad-systems. In section 4, we present context acyclic quad-systems and its properties. We give an account of relevant related works in section 5, and conclude in section 6.

2 Contextualized Quad-Systems

In this section, we formalize the notion of a quad-system and its semantics. For any vector or sequence $\boldsymbol{x}$, we denote by $\|\boldsymbol{x}\|$ the number of symbols in $\boldsymbol{x}$, and by $\{\boldsymbol{x}\}$ the set of symbols in $\boldsymbol{x}$. For any sets A and B, $A \to B$ denotes the set of all functions from set A to set B. Given the set of URIs $\mathbf{U}$, the set of blank nodes $\mathbf{B}$, and the set of literals $\mathbf{L}$, the set $\mathbf{C} = \mathbf{U} \uplus \mathbf{B} \uplus \mathbf{L}$ are called the set of (RDF) constants. Any $(s, p, o) \in \mathbf{C} \times \mathbf{C} \times \mathbf{C}$ is called a generalized RDF triple (from now on, just triple). A graph is defined as a set of triples. A *Quad* is a tuple of the form $c\colon (s, p, o)$, where (s, p, o) is a triple and c is a URI[4], called the *context identifier* that denotes the context of the RDF triple. A *quad-graph* is defined as a set of quads. For any quad-graph Q and any context identifier c, we denote by $graph_Q(c)$ the set $\{(s, p, o)|c\colon (s, p, o) \in Q\}$. We denote by $Q_{\mathcal{C}}$ the quad-graph whose set of context identifiers is $\mathcal{C}$. Let $\mathbf{V}$ be the set of variables, any element of the set $\mathbf{C}^{\mathbf{V}} = \mathbf{V} \cup \mathbf{C}$ is a *term*. Any $(s, p, o) \in \mathbf{C}^{\mathbf{V}} \times \mathbf{C}^{\mathbf{V}} \times \mathbf{C}^{\mathbf{V}}$ is called a *triple pattern*, and an expression of the form $c\colon (s, p, o)$, where (s, p, o) is a triple pattern, c a context identifier, is called a *quad pattern*. A triple pattern t, whose variables are elements of the vector $\boldsymbol{x}$ or elements of the vector $\boldsymbol{y}$ is written as $t(\boldsymbol{x}, \boldsymbol{y})$. For any function $f\colon A \to B$, the *restriction* of f to a set A', is the mapping $f|_{A'}$ from $A' \cap A$ to B s.t. $f|_{A'}(a) = f(a)$, for each $a \in A \cap A'$. For any triple pattern $t = (s, p, o)$ and a function μ from $\mathbf{V}$ to a set A, $t[\mu]$ denotes $(\mu'(s), \mu'(p), \mu'(o))$, where μ' is an extension of μ to $\mathbf{C}$ s.t. $\mu'|_{\mathbf{C}}$ is the identity function. For any set of triple patterns G, $G[\mu]$ denotes $\bigcup_{t \in G} t[\mu]$. For any vector of constants $\boldsymbol{a} = \langle a_1, \ldots, a_{\|\boldsymbol{a}\|} \rangle$, and vector of variables $\boldsymbol{x}$ of the same length, $\boldsymbol{x}/\boldsymbol{a}$ is the function μ s.t. $\mu(x_i) = a_i$, for $1 \leq i \leq \|\boldsymbol{a}\|$. We use the notation $t(\boldsymbol{a}, \boldsymbol{y})$ to denote $t(\boldsymbol{x}, \boldsymbol{y})[\boldsymbol{x}/\boldsymbol{a}]$.

Bridge rules (BRs) *Bridge rules (BR)* enables knowledge propagation across contexts. Formally, a BR is an expression of the form:

$$\forall \boldsymbol{x} \forall \boldsymbol{z}\ [c_1\colon t_1(\boldsymbol{x}, \boldsymbol{z}) \wedge \ldots \wedge c_n\colon t_n(\boldsymbol{x}, \boldsymbol{z}) \to \exists \boldsymbol{y}\ c'_1\colon t'_1(\boldsymbol{x}, \boldsymbol{y}) \wedge \ldots \wedge c'_m\colon t'_m(\boldsymbol{x}, \boldsymbol{y})] \quad (1)$$

where $c_1, \ldots, c_n, c'_1, \ldots, c'_m$ are context identifiers, $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}$ are vectors of variables s.t. $\{\boldsymbol{x}\}, \{\boldsymbol{y}\}$, and $\{\boldsymbol{z}\}$ are pairwise disjoint. $t_1(\boldsymbol{x}, \boldsymbol{z})$, ..., $t_n(\boldsymbol{x}, \boldsymbol{z})$ are triple patterns which do not contain blank-nodes, and whose set of variables are from $\boldsymbol{x}$ or $\boldsymbol{z}$. $t'_1(\boldsymbol{x}, \boldsymbol{y})$, ...,$t'_m(\boldsymbol{x}, \boldsymbol{y})$ are triple patterns, whose set of variables are from $\boldsymbol{x}$ or $\boldsymbol{y}$, and also does not contain blank-nodes. For any BR, r, of the form (1), $body(r)$ is the set of quad patterns $\{c_1\colon t_1(\boldsymbol{x}, \boldsymbol{z}), \ldots, c_n\colon t_n(\boldsymbol{x}, \boldsymbol{z})\}$, and $head(r)$ is the set of quad patterns $\{c'_1\colon t'_1(\boldsymbol{x}, \boldsymbol{y}), \ldots c'_m\colon t'_m(\boldsymbol{x}, \boldsymbol{y})\}$.

Definition 1 (Quad-System). *A* quad-system $QS_{\mathcal{C}}$ *is defined as a pair* $\langle Q_{\mathcal{C}}, R \rangle$*, where* $Q_{\mathcal{C}}$ *is a quad-graph, whose set of context identifiers is* $\mathcal{C}$*, and* R *is a set of BRs.*

[4] Although, in general a context identifier can be a constant, for the ease of notation, we restrict them to be a URI

For any quad-graph $Q_{\mathcal{C}}$ (BR r), its symbols size $\|Q_{\mathcal{C}}\|$ ($\|r\|$) is the number of symbols required to print $Q_{\mathcal{C}}$ (r). Hence, $\|Q_{\mathcal{C}}\| \approx 4*|Q_{\mathcal{C}}|$, where $|Q_{\mathcal{C}}|$ denotes the cardinality of the set $Q_{\mathcal{C}}$. Note that $|Q_{\mathcal{C}}|$ equals the number of quads in $Q_{\mathcal{C}}$. For a BR r, $\|r\| \approx 4 * k$, where k is the number of quad-patterns in r. For a set of BRs R, its size $\|R\|$ is given as $\Sigma_{r \in R}\|r\|$. For any quad-system $QS_{\mathcal{C}} = \langle Q_{\mathcal{C}}, R\rangle$, its size $\|QS_{\mathcal{C}}\| = \|Q_{\mathcal{C}}\| + \|R\|$.

Semantics In order to provide a semantics for enabling reasoning over a quad-system, we need to use a local semantics for each context to interpret the knowledge pertaining to it. Since the primary goal of this paper is a decision procedure for query answering over quad-systems based on forward chaining, we consider the following desiderata for the choice of the local semantics:

- there exists a set of inference rules and an operation lclosure() that computes the deductive closure of a graph w.r.t to the local semantics using the inference rules.
- given a finite graph as input, the lclosure() operation, terminates with a finite graph as output in polynomial time whose size is polynomial w.r.t. to the input set.

Some of the alternatives for the local semantics satisfying the above mentioned criterion are Simple, RDF, RDFS [19], OWL-Horst [20] etc. Assuming that a local semantics has been fixed, for any context c, we denote by $I^c = \langle \Delta^c, \cdot^c \rangle$ an interpretation structure for the local semantics, where Δ^c is the interpretation domain, $\cdot^c$ the corresponding interpretation function. Also $\models_{\text{local}}$ denotes the local satisfaction relation between a local interpretation structure and a graph. Given a quad graph $Q_{\mathcal{C}}$, a *distributed interpretation structure* is an indexed set $\mathcal{I}^{\mathcal{C}} = \{I^c\}_{c \in \mathcal{C}}$, where I^c is a local interpretation structure, for each $c \in \mathcal{C}$. We define the satisfaction relation $\models$ between a distributed interpretation structure $\mathcal{I}^{\mathcal{C}}$ and a quad-system $QS_{\mathcal{C}}$ as:

Definition 2 (Model of a Quad-System). *A distributed interpretation structure $\mathcal{I}^{\mathcal{C}} = \{I^c\}_{c \in \mathcal{C}}$ satisfies a quad-system $QS_{\mathcal{C}} = \langle Q_{\mathcal{C}}, R\rangle$, in symbols $\mathcal{I}^{\mathcal{C}} \models QS_{\mathcal{C}}$, iff all the following conditions are satisfied:*

1. *$I^c \models_{local} graph_{Q_{\mathcal{C}}}(c)$, for each $c \in \mathcal{C}$;*
2. *$a^{c_i} = a^{c_j}$, for any $a \in \mathbf{C}$, $c_i, c_j \in \mathcal{C}$;*
3. *for each BR $r \in R$ of the form (1) and for each $\sigma \in \mathbf{V} \rightarrow \Delta^{\mathcal{C}}$, where $\Delta^{\mathcal{C}} = \bigcup_{c \in \mathcal{C}} \Delta^c$, if*
$$I^{c_1} \models_{local} t_1(\boldsymbol{x}, \boldsymbol{z})[\sigma], ..., I^{c_n} \models_{local} t_n(\boldsymbol{x}, \boldsymbol{z})[\sigma],$$
then there exists function $\sigma' \supseteq \sigma$, s.t.
$$I^{c'_1} \models_{local} t'_1(\boldsymbol{x}, \boldsymbol{y})[\sigma'], ..., I^{c'_m} \models_{local} t'_m(\boldsymbol{x}, \boldsymbol{y})[\sigma'].$$

Condition 1 in the above definition ensures that for any model $\mathcal{I}^{\mathcal{C}}$ of a quad-graph, each $I^c \in \mathcal{I}^{\mathcal{C}}$ is a local model of the set of triples in context c. Condition 2 ensures that any constant c is rigid, i.e. represents the same resource across a quad-graph, irrespective of the context in which it occurs. Condition 3 ensure that any model of a quad-system satisfies each BR in it. Any $\mathcal{I}^{\mathcal{C}}$ s.t. $\mathcal{I}^{\mathcal{C}} \models QS_{\mathcal{C}}$ is said to be a model of $QS_{\mathcal{C}}$. A quad-system $QS_{\mathcal{C}}$ is said to be *consistent* if there exists a model $\mathcal{I}^{\mathcal{C}}$, s.t. $\mathcal{I}^{\mathcal{C}} \models QS_{\mathcal{C}}$, and otherwise said to be *inconsistent*. For any quad-system $QS_{\mathcal{C}} = \langle Q_{\mathcal{C}}, R\rangle$, it can be the case that $graph_{Q_{\mathcal{C}}}(c)$ is locally consistent, for each $c \in \mathcal{C}$, whereas $QS_{\mathcal{C}}$ is not consistent. This is because the set of BRs R adds more knowledge to the quad-system, and restricts the set of models that satisfy the quad-system.

Definition 3 (Quad-system entailment). *(a) A quad-system $QS_{\mathcal{C}}$ entails a quad $c\colon (s, p, o)$, in symbols $QS_{\mathcal{C}} \models c\colon (s, p, o)$, iff for any distributed interpretation structure $\mathcal{I}^{\mathcal{C}}$, if $\mathcal{I}^{\mathcal{C}} \models QS_{\mathcal{C}}$ then $\mathcal{I}^{\mathcal{C}} \models \langle\{c\colon (s, p, o)\}, \emptyset\rangle$. (b) A quad-system $QS_{\mathcal{C}}$ entails a quad-graph $Q'_{\mathcal{C}'}$, in symbols $QS_{\mathcal{C}} \models Q'_{\mathcal{C}'}$ iff $QS_{\mathcal{C}} \models c\colon (s, p, o)$ for any $c\colon (s, p, o) \in Q'_{\mathcal{C}'}$. (c) A quad-system $QS_{\mathcal{C}}$ entails a BR r iff for any $\mathcal{I}^{\mathcal{C}}$, if $\mathcal{I}^{\mathcal{C}} \models QS_{\mathcal{C}}$ then $\mathcal{I}^{\mathcal{C}} \models \langle\emptyset, \{r\}\rangle$. (d) For a set of BRs R, $QS_{\mathcal{C}} \models R$ iff $QS_{\mathcal{C}} \models r$, for every $r \in R$. (e) Finally, a quad-system $QS_{\mathcal{C}}$ entails another quad-system $QS'_{\mathcal{C}'} = \langle Q'_{\mathcal{C}'}, R'\rangle$, in symbols $QS_{\mathcal{C}} \models QS'_{\mathcal{C}'}$ iff $QS_{\mathcal{C}} \models Q'_{\mathcal{C}'}$ and $QS_{\mathcal{C}} \models R'$.*

We call the decision problems (DPs) corresponding to the entailment problems (EPs) in (a), (b), (c), (d), and (e) as *quad EP, quad-graph EP, BR EP, BRs EP, and quad-system EP*, respectively.

3 Query Answering on Quad-Systems

In the realm of quad-systems, the classical conjunctive queries or select-project-join queries are slightly extended to what we call *Contextualized Conjunctive Queries* (CCQs). A CCQ $CQ(\boldsymbol{x})$ is an expression of the form:

$$\exists \boldsymbol{y}\ q_1(\boldsymbol{x}, \boldsymbol{y}) \wedge ... \wedge q_p(\boldsymbol{x}, \boldsymbol{y}) \tag{2}$$

where q_i, for $i = 1, ..., p$ are quad patterns over vectors of *free variables* $\boldsymbol{x}$ and *quantified variables* $\boldsymbol{y}$. A CCQ is called a boolean CCQ if it does not have any free variables. For any CCQ $CQ(\boldsymbol{x})$ and a vector $\boldsymbol{a}$ of constants s.t. $\|\boldsymbol{x}\| = \|\boldsymbol{a}\|$, $CQ(\boldsymbol{a})$ is boolean. A vector $\boldsymbol{a}$ is an *answer* for a CCQ $CQ(\boldsymbol{x})$ w.r.t. structure $\mathcal{I}_{\mathcal{C}}$, in symbols $\mathcal{I}_{\mathcal{C}} \models CQ(\boldsymbol{a})$, iff there exists assignment $\mu\colon \{\boldsymbol{y}\} \rightarrow \mathbf{B}$ s.t. $\mathcal{I}_{\mathcal{C}} \models \bigcup_{i=1,...,p} q_i(\boldsymbol{a}, \boldsymbol{y})[\mu]$. A vector $\boldsymbol{a}$ is a *certain answer* for a CCQ $CQ(\boldsymbol{x})$ over a quad-system $QS_{\mathcal{C}}$, iff $\mathcal{I}_{\mathcal{C}} \models CQ(\boldsymbol{a})$, for every model $\mathcal{I}_{\mathcal{C}}$ of $QS_{\mathcal{C}}$. Given a quad-system $QS_{\mathcal{C}}$, a CCQ $CQ(\boldsymbol{x})$, and a vector $\boldsymbol{a}$, DP of determining whether $QS_{\mathcal{C}} \models CQ(\boldsymbol{a})$ is called the *CCQ EP*. It can be noted that the other DPs over quad-systems that we have seen are reducible to CCQ EP. Hence, in this paper, we primarily focus on the CCQ EP.

dChase of a Quad-System In order to do query answering over a quad-system, we employ what has been called in the literature, a *chase* [22, 1], specifically, we adopt notion of the *skolem chase* given in Marnette [26] and Cuenca Grau et al [9]. In order to fit the framework of quad-systems, we extend the standard notion of chase to a *distributed chase*, abbreviated *dChase*. In the following, we show how the dChase of a quad-system can be constructed.

For any BR r of the form (1), the *skolemization* $sk(r)$ is the result of replacing each $y_i \in \{\boldsymbol{y}\}$ with a globally unique Skolem function f_i^r, s.t. $f_i^r\colon\ \mathbf{C}^{\|\boldsymbol{x}\|} \rightarrow \mathbf{B}_{sk}$, where $\mathbf{B}_{sk}$ is a fresh set of blank nodes called *skolem blank nodes*. Intuitively, for every distinct vector $\boldsymbol{a}$ of constants, with $\|\boldsymbol{a}\| = \|\boldsymbol{x}\|$, $f_i^r(\boldsymbol{a})$ is a fresh blank node, whose node id is a hash of $\boldsymbol{a}$. Let $\boldsymbol{f}^r = \langle f_1^r, ..., f_{\|\boldsymbol{y}\|}^r\rangle$ be a vector of distinct Skolem functions; for any BR r the form (1), with slight abuse (Datalog notation) we write its skolemization $sk(r)$ as follows:

$$c_1 : t_1(\boldsymbol{x}, \boldsymbol{z}), ..., c_n\colon t_n(\boldsymbol{x}, \boldsymbol{z}) \rightarrow c'_1\colon t'_1(\boldsymbol{x}, \boldsymbol{f}^r), ..., c'_m\colon t'_m(\boldsymbol{x}, \boldsymbol{f}^r) \tag{3}$$

Moreover, a skolemized BR r of the form (3) can be replaced by the following semantically equivalent set of formulas, whose symbol size is worst case quadratic w.r.t $\|r\|$:

$$\begin{aligned}&\{c_1 : t_1(\boldsymbol{x}, \boldsymbol{z}), ..., c_n : t_n(\boldsymbol{x}, \boldsymbol{z}) \rightarrow c'_1 : t'_1(\boldsymbol{x}, \boldsymbol{f}^r), \\ &..., \\ &c_1 : t_1(\boldsymbol{x}, \boldsymbol{z}), ..., c_n : t_n(\boldsymbol{x}, \boldsymbol{z}) \rightarrow c'_m : t'_m(\boldsymbol{x}, \boldsymbol{f}^r)\}\end{aligned} \tag{4}$$

Note that each BR in the above set has exactly one quad pattern with optional function symbols in its head part. Also note that a BR with out function symbols can be replaced with a set of BRs with single quad-pattern heads. Hence, w.l.o.g, we assume that any BR in a skolemized set $sk(R)$ of BRs is of the form (4). For any quad-graph $Q_{\mathcal{C}}$ and a skolemized BR r of the form (4), *application* of r on $Q_{\mathcal{C}}$, denoted by $r(Q_{\mathcal{C}})$, is given as:

$$r(Q_{\mathcal{C}}) = \bigcup_{\mu \in \mathbf{V} \rightarrow \mathbf{C}} \{\, c'_1 : t'_1(\boldsymbol{x}, \boldsymbol{f}^r)[\mu] \mid c_1 : t_1(\boldsymbol{x}, \boldsymbol{z})[\mu] \in Q_{\mathcal{C}}, ..., c_n : t_n(\boldsymbol{x}, \boldsymbol{z})[\mu] \in Q_{\mathcal{C}} \,\}$$

For any set of skolemized BRs R, application of R on $Q_{\mathcal{C}}$ is given by: $R(Q_{\mathcal{C}}) = \bigcup_{r \in R} r(Q_{\mathcal{C}})$. For any quad-graph $Q_{\mathcal{C}}$, we define:

$$\mathsf{lclosure}(Q_{\mathcal{C}}) = \bigcup_{c \in \mathcal{C}} \{c : (s, p, o) \mid (s, p, o) \in \mathsf{lclosure}(graph_{Q_{\mathcal{C}}}(c))\}$$

For any quad-system $QS_{\mathcal{C}} = \langle Q_{\mathcal{C}}, R \rangle$, *generating BRs* R_F is the set of BRs in $sk(R)$ with function symbols, and the *non-generating BRs* is the set $R_I = sk(R) \setminus R_F$.
Let $dChase_0(QS_{\mathcal{C}}) = \mathsf{lclosure}(Q_{\mathcal{C}})$; for $i \in \mathbb{N}$, $dChase_{i+1}(QS_{\mathcal{C}}) =$

$$\begin{cases} \mathsf{lclosure}(dChase_i(QS_{\mathcal{C}}) \cup R_I(dChase_i(QS_{\mathcal{C}}))), & \text{if } R_I(dChase_i(QS_{\mathcal{C}})) \not\subseteq dChase_i(QS_{\mathcal{C}}); \\ \mathsf{lclosure}(dChase_i(QS_{\mathcal{C}}) \cup R_F(dChase_i(QS_{\mathcal{C}}))), & \text{otherwise;} \end{cases}$$

The dChase of $QS_{\mathcal{C}}$, denoted $dChase(QS_{\mathcal{C}})$, is given as:

$$dChase(QS_{\mathcal{C}}) = \bigcup_{i \in \mathbb{N}} dChase_i(QS_{\mathcal{C}})$$

Intuitively, $dChase_i(QS_{\mathcal{C}})$ can be thought of as the state of $dChase(QS_{\mathcal{C}})$ at the end of iteration i. It can be noted that, if there exists i s.t. $dChase_i(QS_{\mathcal{C}}) = dChase_{i+1}(QS_{\mathcal{C}})$, then $dChase(QS_{\mathcal{C}}) = dChase_i(QS_{\mathcal{C}})$. An iteration i, s.t. $dChase_i(QS_{\mathcal{C}})$ is computed by the application of the set of (resp. non-)generating BRs R_F (resp. R_I), on $dChase_{i-1}(QS_{\mathcal{C}})$ is called a (resp. non-)*generating iteration*. The dChase $dChase(QS_{\mathcal{C}})$ of a consistent quad-system $QS_{\mathcal{C}}$ is a *universal model* [10] of the quad-system, i.e. it is a model of $QS_{\mathcal{C}}$, and for any model $\mathcal{I}_{\mathcal{C}}$ of $QS_{\mathcal{C}}$, there is a homomorphism from $dChase(QS_{\mathcal{C}})$ to $\mathcal{I}_{\mathcal{C}}$. Hence, for any boolean CCQ $CQ()$, $QS_{\mathcal{C}} \models CQ()$ iff there exists a map $\mu\colon \mathbf{V}(CQ) \rightarrow \mathbf{C}$ s.t. $\{CQ()\}[\mu] \subseteq dChase(QS_{\mathcal{C}})$. We call the sequence $dChase_0(QS_{\mathcal{C}})$, $dChase_1(QS_{\mathcal{C}}), ...$, the *dChase sequence* of $QS_{\mathcal{C}}$. The following lemma shows that in a dChase sequence of a quad-system, the result of a single generating iteration and a subsequent number of non-generating iterations causes only an exponential blow up in size.

Lemma 1. *For a quad-system $QS_{\mathcal{C}} = \langle Q_{\mathcal{C}}, R\rangle$, the following holds: (i) if $i \in \mathbb{N}$ is a generating iteration, then $\|dChase_i(QS_{\mathcal{C}})\| = \mathcal{O}(\|dChase_{i-1}(QS_{\mathcal{C}})\|^{\|R\|})$, (ii) suppose $i \in \mathbb{N}$ is a generating iteration, and for any $j \geq 1$, $i+1$, ..., $i+j$ are non-generating iterations, then $\|dChase_{i+j}(QS_{\mathcal{C}})\| = \mathcal{O}(\|dChase_{i-1}(QS_{\mathcal{C}})\|^{\|R\|})$, (iii) for any iteration k, $dChase_k(QS_{\mathcal{C}})$ can be computed in time $\mathcal{O}(\|dChase_{k-1}(QS_{\mathcal{C}})\|^{\|R\|})$.*

Proof (sketch). (i) R can be applied on $dChase_{i-1}(QS_{\mathcal{C}})$ by grounding R to the set of constants in $dChase_{i-1}(QS_{\mathcal{C}})$, the number of such groundings is of the order $\mathcal{O}(\|dChase_{i-1}(QS_{\mathcal{C}})\|^{\|R\|})$, $\|R(dChase_{i-1}(QS_{\mathcal{C}}))\| = \mathcal{O}(\|R\| * \|dChase_{i-1}(QS_{\mathcal{C}})\|^{\|R\|})$. Since lclosure only increases the size polynomially, $\|dChase_i(QS_{\mathcal{C}})\| = \mathcal{O}(\|dChase_{i-1}(QS_{\mathcal{C}})\|^{\|R\|})$.

(ii) From (i) we know that $\|R(dChase_{i-1}(QS_{\mathcal{C}}))\| = \mathcal{O}(\|dChase_{i-1}(QS_{\mathcal{C}})\|^{\|R\|})$. Since, no new constant is introduced in any subsequent non-generating iterations, and since any quad contains only four constants, the set of constants in any subsequent dChase iteration is $\mathcal{O}(4 * \|dChase_{i-1}(QS_{\mathcal{C}})\|^{\|R\|})$. Since only these many constants can appear in positions c, s, p, o of any quad generated in the subsequent iterations, the size of $dChase_{i+j}(QS_{\mathcal{C}})$ can only increase polynomially, which means that $\|dChase_{i+j}(QS_{\mathcal{C}})\| = \mathcal{O}(\|dChase_{i-1}(QS_{\mathcal{C}})\|^{\|R\|})$.

(iii) Since any dChase iteration k involves the following two operations: (a) lclosure(), and (b) computing $R(dChase_{k-1}(QS_{\mathcal{C}}))$. (a) can be done in PTIME w.r.t to its input. (b) can be done in the following manner: ground R to the set of constants in $dChase_{i-1}(QS_{\mathcal{C}})$; then for each grounding g, if $body(g) \subseteq dChase_{i-1}(QS_{\mathcal{C}})$, then add $head(g)$ to $R(dChase_{k-1}(QS_{\mathcal{C}}))$. Since, the number of such groundings is of the order $\mathcal{O}(\|dChase_{k-1}(QS_{\mathcal{C}})\|^{\|R\|})$, and checking if each grounding is contained in $dChase_{k-1}(QS_{\mathcal{C}})$, can be done in time polynomial in $\|dChase_{k-1}(QS_{\mathcal{C}})\|$, the time taken for (b) is $\mathcal{O}(\|dChase_{k-1}(QS_{\mathcal{C}})\|^{\|R\|})$. Hence, any iteration k can be done in time $\mathcal{O}(\|dChase_{k-1}(QS_{\mathcal{C}})\|^{\|R\|})$. □

Although, we now know how to compute the dChase of a quad-system, which can be used for deciding CCQ EP, it turns out that for the class of quad-systems whose BRs are of the form (1), which we call *unrestricted quad-systems*, the dChase can be infinite. This raises the question if there are other approaches that can be used, for instance similar problem arises in DLs with value creation, due to the presence of existential quantifiers, whereas the approaches like the one in Glim et al. [16] provides an algorithm for CQ entailment based on query rewriting.

Theorem 1. *The CCQ EP over unrestricted quad-systems is undecidable.*

Proof (sketch). We show that the well known undecidable problem of non-emptiness of intersection of context-free grammars (CFGs) is reducible to the CCQ entailment problem. Given two CFGs, $G_1 = \langle V_1, T, S_1, P_1\rangle$ and $G_2 = \langle V_2, T, S_2, P_2\rangle$, where V_1, V_2 are the set of variables, T s.t. $T \cap (V_1 \cup V_2) = \emptyset$ is the set of terminals. $S_1 \in V_1$ is the start symbol of G_1, and P_1 are the set of PRs of the form $v \rightarrow \boldsymbol{w}$, where $v \in V$, $\boldsymbol{w}$ is a sequence of the form $w_1 ... w_n$, where $w_i \in V_1 \cup T$. Similarly s_2, P_2 is defined. Deciding whether the language generated by the grammars $L(G_1)$ and $L(G_2)$ have non-empty intersection is known to be undecidable [18].

Given two CFGs $G_1 = \langle V_1, T, S_1, P_1 \rangle$ and $G_2 = \langle V_2, T, S_2, P_2 \rangle$, we encode grammars G_1, G_2 into a quad-system $QS_c = \langle Q_c, R \rangle$, with only a single context identifier c. Each PR $r = v \rightarrow \boldsymbol{w} \in P_1 \cup P_2$, with $\boldsymbol{w} = w_1 w_2 w_3 .. w_n$, is encoded as a BR of the form: $c\colon (x_1, w_1, x_2), c\colon (x_2, w_2, x_3), ..., c\colon (x_n, w_n, x_{n+1}) \rightarrow c\colon (x_1, v, x_{n+1})$, where $x_1, .., x_{n+1}$ are variables. For each terminal symbol $t_i \in T$, R contains a BR of the form: $c\colon (x, \texttt{rdf:type}, C) \rightarrow \exists y\ c\colon (x, t_i, y), c\colon (y, \texttt{rdf:type}, C)$ and Q_c is the singleton: $\{c\colon (a, \texttt{rdf:type}, C)\}$. It can be observed that:

$$QS_c \models \exists y\ c\colon (a, S_1, y)\ \wedge\ c\colon (a, S_2, y) \Leftrightarrow L(G_1) \cap L(G_2) \neq \emptyset$$

We refer the reader to Appendix for the complete proof. □

4 Context Acyclic Quad-Systems: A decidable class

In the previous section, we saw that query answering on unrestricted quad-systems is undecidable, in general. We in the following define a class of quad-systems for which query entailment is decidable. The class has the property that algorithms based on forward chaining, for deciding query entailment, can straightforwardly be implemented (by minor extensions) on existing quad stores. It should be noted that the technique we propose is reminiscent of the *Weak acyclicity* [13, 11] technique used in the realm of Datalog+-.

Consider a BR r of the form: $c_1 : t_1(\boldsymbol{x}, \boldsymbol{z}), c_2 : t_2(\boldsymbol{x}, \boldsymbol{z}) \rightarrow \exists \boldsymbol{y}\ c_3 : t_3(\boldsymbol{x}, \boldsymbol{y}), c_4 : t_4(\boldsymbol{x}, \boldsymbol{y})$. Since such a rule triggers propagation of knowledge in a quad-system, specifically triples from the source contexts c_1, c_2 to the target contexts c_3, c_4 in a quad-system. As shown in Fig. 1, we can view a BR as a propagation rule across distinct compartments of knowledge, divided as contexts. For any BR of the form (1), each context in the set $\{c'_1, ..., c'_m\}$ is said to depend on the set of contexts $\{c_1, ..., c_n\}$. In a quad-system $QS_{\mathcal{C}} = \langle Q_{\mathcal{C}}, R \rangle$, for any $r \in R$, of the form (1), any context whose identifier is in the set $\{c \mid c\colon (s, p, o) \in head(r), s$ or p or o is an existentially quantified variable$\}$, is called a *triple generating context* (TGC). One can analyze the set of BRs in a quad-system $QS_{\mathcal{C}}$ using a context dependency graph, which is a directed graph, whose nodes are context identifiers in $\mathcal{C}$, s.t. the nodes corresponding to TGCs are marked with a $*$, and whose edges are constructed as follows: for each BR of the form (1), there exists an edge from each c_i to c'_j, for $i = 1, ..., n, j = 1, ..., m$. A quad-system is said to be *context acyclic*, iff its context dependency graph does not contain cycles involving TGCs.

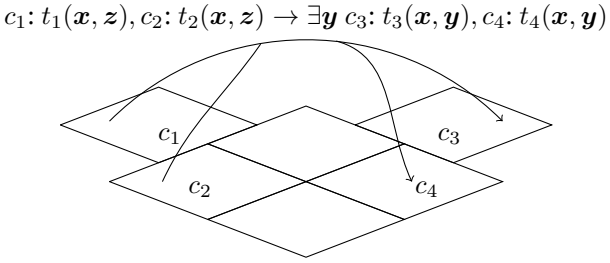

Fig. 1

Example 1. Consider a quad-system, whose set of BRs R are:

$$c_1\colon (x_1, x_2, \mathsf{U}_1) \rightarrow \exists y_1\ c_2\colon (x_1, x_2, y_1), c_3\colon (x_2, \texttt{rdf:type}, \texttt{rdf:Property}) \quad (5)$$

$$c_2\colon (x_1, x_2, z_1) \rightarrow c_1\colon (x_1, x_2, \mathsf{U}_1) \quad (6)$$

$$c_3\colon (x_1, x_2, x_3) \rightarrow c_1\colon (x_1, x_2, x_3)$$

where U_1 be a URI, whose corresponding dependency graph is shown in Fig. 2. Note that the node corresponding to the triple generating context c_2 is marked with a '$*$' symbol. Since the cycle (c_1, c_2, c_1) in the quad-system contains c_2 which is a TGC, the quad-system is not context acyclic.

In a context acyclic quad-system $QS_{\mathcal{C}}$, since there exists no cyclic path through any TGC node in the context dependency graph, there exists a set of TGCs $\mathcal{C}' \subseteq \mathcal{C}$ s.t. for any $c \in \mathcal{C}'$, there exists no incoming path[5] from a TGC to c. We call such TGCs, *level-1 TGCs*. In other words, a TGC c is a level-1 TGC, if for any $c' \in \mathcal{C}$, there exists an incoming path from c' to c, implies c' is not a TGC. For $l \geq 1$, a level-l+1 TGC c is a TGC that has an incoming path from a level-l TGC, and for any incoming path from a level-l' TGC to c, is s.t. $l' \leq l$. Extending the notion of level also to the non-TGCs, we say that any non-TGC that does not have any incoming paths from a TGC is at level-0; we say that any non-TGC $c \in \mathcal{C}$ is at level-l, if there exists an incoming path from a level-l TGC to c, and for any incoming path from a level-l' TGC to c, is s.t. $l' \leq l$. Hence, the set of contexts in a context acyclic quad-system can be partitioned using the above notion of levels.

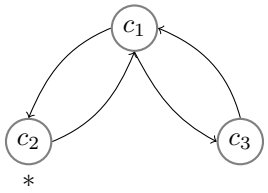

Fig. 2: Context Dependency graph

Definition 4. *For a quad-system $QS_{\mathcal{C}}$, a context $c \in \mathcal{C}$ is said to be* saturated *in an iteration i, iff for any quad of the form $c\colon (s, p, o)$, $c\colon (s, p, o) \in dChase(QS_{\mathcal{C}})$ implies $c\colon (s, p, o) \in dChase_i(QS_{\mathcal{C}})$.*

Intuitively, context c is saturated in the dChase iteration i, if no new quad of the form $c\colon (s, p, o)$ will be generated in any $dChase_k(QS_{\mathcal{C}})$, for any $k > i$. The following lemma gives the relation between the saturation of a context and the required number of dChase iterations, for a context acyclic quad-system.

Lemma 2. *For any context acyclic quad-system, the following holds: (i) any level-0 context is saturated before the first generating iteration, (ii) any level-1 TGC is saturated after the first generating iteration, (iii) any level-k context is saturated before the $k+1$th generating iteration.*

Proof. Let $QS_{\mathcal{C}} = \langle Q_{\mathcal{C}}, R \rangle$ be the quad-system, whose first generating iteration is i.

(i) for any level-0 context c, any BR $r \in R$, and any quad-pattern of the form $c\colon (s, p, o)$, if $c\colon (s, p, o) \in head(r)$, then for any c' s.t. $c'\colon (s', p', o')$ occurs in $body(r)$

[5] assume that paths have at least one edge

implies that c' is a level-0 context and r is a non-generating BR. Also, since c' is a level-0 context, the same applies to c'. Hence, it turns out that only non-generating BRs can bring triples to any level-0 context. Since at the end of iteration $i-1$, $dChase_{i-1}(QS_{\mathcal{C}})$ is closed w.r.t. the set of non-generating BRs (otherwise, by construction of dChase, i would not be a generating iteration). This implies that c is saturated before the first generating iteration i.

(ii) for any level-1 TGC c, any BR $r \in R$, and any quad-pattern $c\colon (s, p, o)$, if $c\colon (s, p, o) \in head(r)$, then for any c' s.t. $c'\colon (s', p', o')$ occurs in $body(r)$ implies that c' is a level-0 context (Otherwise level of c would be greater than 1). This means that only contexts from which triples get propagated to c are level-0 contexts. From (i) we know that all the level-0 contexts are saturated before ith iteration, and since during the ith iteration R_F is applied followed by the lclosure() operation (R_I need not be applied, since $dChase_{i-1}(QS_{\mathcal{C}})$ is closed w.r.t. R_I), c is saturated after iteration i, the 1st generating iteration.

(iii) can be obtained from generalization of (i) and (ii), and from the fact that any level-k context can only have incoming paths from contexts whose levels are less than or equal to k. □

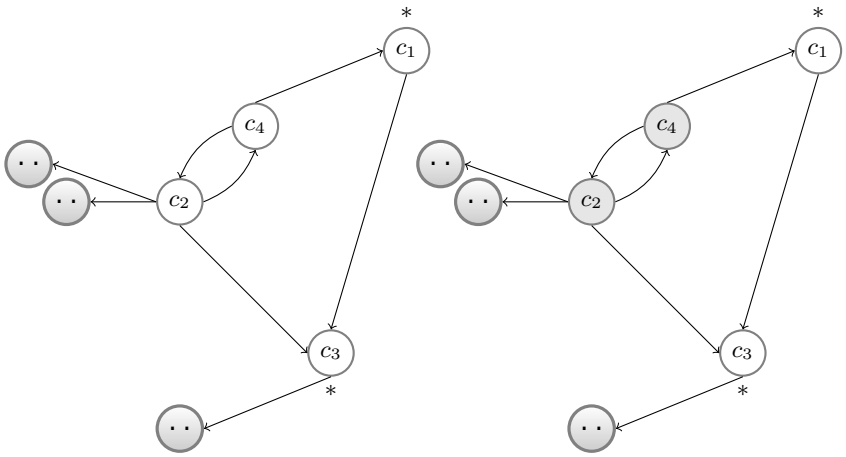

(a) (b)

Fig. 3

Example 2. Consider the dependency graph in Fig. 3a, where .. indicates part of the graph that is not under the scope of our discussion. The TGCs nodes c_1 and c_3 are marked with a $*$. It can be seen that both c_2 and c_4 are level-0 contexts, since they do not have any incoming paths from TGCs. Since the only incoming paths to context c_1 are from c_2 and c_4, which are not TGCs, c_1 is a level-1 TGC. Context c_3 is a level-2 TGC, since it has an incoming path from the level-1 TGC c_1, and has no incoming path

from a TGC whose level is greater than 1. Since the level-0 contexts only have incoming paths from level-0 contexts and only appear on the head part of non-generating BRs, before first generating iteration, all the level-0 TGCs becomes saturated, as the set of non-generating BRs R_I has been exhaustively applied. This situation is reflected in Fig. 3b, where the saturated nodes are shaded with gray. Note that after the first and second generating iterations c_1 and c_3 also become saturated, respectively.

The following lemma shows that for context acyclic quad-systems, there exists a finite bound on the size and computation time of its dChase.

Lemma 3. *For any context acyclic quad-system $QS_{\mathcal{C}} = \langle Q_{\mathcal{C}}, R\rangle$, the following holds: (i) the number of dChase iterations is finite, (ii) size of the dChase $\|dChase(QS_{\mathcal{C}})\| = \mathcal{O}(2^{2^{\|QS_{\mathcal{C}}\|}})$, (iii) computing $dChase(QS_{\mathcal{C}})$ is in* 2EXPTIME*, (iv) if R and the set of schema triples in $Q_{\mathcal{C}}$ is fixed, then $\|dChase(QS_{\mathcal{C}})\|$ is a polynomial in $\|QS_{\mathcal{C}}\|$, and computing $dChase(QS_{\mathcal{C}})$ is in* PTIME.

Proof. (i) Since $QS_{\mathcal{C}}$ is context-acyclic, all the contexts can be partitioned according to their levels. Also, the number of levels k is s.t. $k \leq |\mathcal{C}|$. Hence, applying lemma 1, before the $k+1$th generating iteration all the contexts becomes saturated, and $k+1$th generating iteration do not produce any new quads, terminating the dChase computation process.

(ii) In the dChase computation process, since by lemma 1, any generating iteration and a sequence of non-generating iterations can only increase the dChase size exponentially in $\|R\|$, the size of the dChase before $k+1$ th generating iteration is $\mathcal{O}(\|dChase_0(QS_{\mathcal{C}})\|^{\|R\|^k})$, which can be written as $\mathcal{O}(\|QS_{\mathcal{C}}\|^{\|R\|^k})$ (†). As seen in (i), there can only be $|\mathcal{C}|$ generating iterations, and a sequence of non-generating iterations. Hence, applying $k = |\mathcal{C}|$ to (†), and taking into account the fact that $|\mathcal{C}| \leq \|QS_{\mathcal{C}}\|$, the size of the dChase $\|dChase(QS_{\mathcal{C}})\| = \mathcal{O}(2^{2^{\|QS_{\mathcal{C}}\|}})$.

(iii) Since in any dChase iteration except the final one, atleast one new quad should be produced and the final dChase can have at most $\mathcal{O}(2^{2^{\|QS_{\mathcal{C}}\|}})$ quads (by ii), the total number of iterations are bounded by $\mathcal{O}(2^{2^{\|QS_{\mathcal{C}}\|}})$ (†). Since from lemma 1, we know that for any iteration i, computing $dChase_i(QS_{\mathcal{C}})$ is of the order $\mathcal{O}(\|dChase_{i-1}(QS_{\mathcal{C}})\|^{\|R\|})$. Since, $\|dChase_{i-1}(QS_{\mathcal{C}})\|$ can at most be $\mathcal{O}(2^{2^{\|QS_{\mathcal{C}}\|}})$, computing $dChase_i(QS_{\mathcal{C}})$ is of the order $\mathcal{O}(2^{\|R\|*2^{\|QS_{\mathcal{C}}\|}})$. Also since $\|R\| \leq \|QS_{\mathcal{C}}\|$, any iteration requires $\mathcal{O}(2^{2^{\|QS_{\mathcal{C}}\|}})$ time (‡). From (†) and (‡), we can conclude that the time required for computing dChase is in 2EXPTIME.

(iv) In (ii) we saw that the size of the dChase before $k+1$th generating iteration is given by $\mathcal{O}(\|QS_{\mathcal{C}}\|^{\|R\|^k})$ ($\diamond$). Since by hypothesis $\|R\|$ is a constant and also the size of the dependency graph and the levels in it. Hence, the expression $\|R\|^k$ in ($\diamond$) amounts to a constant z. Hence, $\|dChase(QS_{\mathcal{C}})\| = \mathcal{O}(\|QS_{\mathcal{C}}\|^z)$. Hence, the size of $dChase(QS_{\mathcal{C}})$ is a polynomial in $\|QS_{\mathcal{C}}\|$.

Also, since in any dChase iteration except the final one, atleast one quad should be produced and the final dChase can have at most $\mathcal{O}(\|QS_{\mathcal{C}}\|^z)$ quads, the total number of iterations are bounded by $\mathcal{O}(\|QS_{\mathcal{C}}\|^z)$ (†). Also from lemma 1, we know that any dChase iteration i, computing $dChase_i(QS_{\mathcal{C}})$ involves two steps: (a) computing $R(dChase_{i-1}(QS_{\mathcal{C}}))$, and (b) computing lclosure(), which can be done in

PTIME in the size of its input. Since computing $R(dChase_{i-1}(QS_{\mathcal{C}}))$ is of the order $\mathcal{O}(\|dChase_{i-1}(QS_{\mathcal{C}})\|^{\|R\|})$, where $|R|$ is a constant and $\|dChase_{i-1}(QS_{\mathcal{C}})\|$ is a polynomial is $\|QS_{\mathcal{C}}\|$, each iteration can be done in time polynomial in $\|QS_{\mathcal{C}}\|$ (‡). From (†) and (‡), it can be concluded that dChase can be computed in PTIME. □

Lemma 4. *For any context acyclic quad-system, the following holds: (i) data complexity of CCQ entailment is in* PTIME *(ii) combined complexity of CCQ entailment is in* 2EXPTIME.

Proof. For a context acyclic quad-system $QS_{\mathcal{C}} = \langle Q_{\mathcal{C}}, R \rangle$, since $dChase(QS_{\mathcal{C}})$ is finite, a boolean CCQ $CQ()$ can naively be evaluated by grounding the set of constants in the chase to the variables in the $CQ()$, and then checking if any of these groundings are contained in $dChase(QS_{\mathcal{C}})$. The number of such groundings can at most be $\|dChase(QS_{\mathcal{C}})\|^{\|CQ()\|}$ (†).

(i) Since for data complexity, the size of the BRs $\|R\|$, the set of schema triples, and $\|CQ()\|$ is fixed to constant. From lemma 3 (iv), we know that under the above mentioned settings the dChase can be computed in PTIME and is polynomial in the size of $QS_{\mathcal{C}}$. Since $\|CQ()\|$ is fixed to a constant, and from (†), binding the set of constants in $dChase(QS_{\mathcal{C}})$ on $CQ()$ still gives a number of bindings that is worst case polynomial in the size of $QS_{\mathcal{C}}$. Since membership of these bindings can checked in the polynomially sized dChase in PTIME, the time required for CCQ evaluation is in PTIME.

(ii) Since in this case $\|dChase(QS_{\mathcal{C}})\| = \mathcal{O}(2^{2^{\|QS_{\mathcal{C}}\|}})$ (‡), from (†) and (‡), binding the set of constants in $\|dChase(QS_{\mathcal{C}}\|$ to variables in $CQ()$ amounts to $\mathcal{O}(2^{\|CQ()\| * 2^{\|QS_{\mathcal{C}}\|}})$ bindings. Since the size of dChase is double exponential in $\|QS_{\mathcal{C}}\|$, checking the membership of each of these bindings can be done in 2EXPTIME. Hence, the combined complexity is in 2EXPTIME. □

Theorem 2. *For any context acyclic quad-system, the following holds: (i) The data complexity of CCQ entailment is* PTIME*-complete, (ii) The combined complexity of CCQ entailment is* 2EXPTIME*-complete.*

For PTIME-hardness of data complexity, it can be shown that the well known problem of 3HornSat, the satisfiability of propositional Horn formulas with atmost 3 literals, and for 2EXPTIME-hardness for the combined complexity, it can be shown that the word problem of a double exponentially time bounded deterministic turing machine, which is a well known 2EXPTIME-hard problem, is reducible to the CCQ entailment problem (see appendix for detailed proof).

Reconsidering the quad-system in example 1, which is not context acyclic. Suppose that the contexts are enabled with RDFS inferencing, i.e $\mathsf{lclosure}() = \mathsf{rdfsclosure}()$. During dChase construction, since any application of rule (5) can only create a triple in c_2 in which the skolem blank node is in the object position, where as the application of rule (6), does not propogate constants in object postion to c_1. Although at a first look, the dChase might seem to terminate, but since the application of the following RDFS inference rule in c_2: $(s, p, o) \rightarrow (o$,`rdf:type`, `rdfs:Resource`), derives a quad of the form c_2: (_:b, `rdf:type`, `rdfs:Resource`), where _:b is the skolem blank-node created by the application of rule (5). Now by application of rule (6) leads

to $c_1 : (_ : b, \texttt{rdf:type}, U_1)$. Since rule (5) is applicable on $c_1 : (_ : b, \texttt{rdf:type}, U_1)$, which again brings a new skolem blank node to c_2, and hence the dChase construction doesn't terminate. Hence, as seen above the notion of context acyclicity can alarm us about such infinite cases.

5 Related Work

Contexts and Distributed Logics The work on contexts began in the 80s when McCarthy [21] proposed context as a solution to the generality problem in AI. After this various studies about logics of contexts mainly in the field of KR was done by Guha [28], *Distributed First Order Logics* by Ghidini et al. [14] and *Local Model Semantics* by Giunchiglia et al. [15]. Primarily in these works contexts are formalized as a first order/propositional theories, and bridge rules were provided to inter-operate the various contexts. Some of the initial works on contexts relevant to semantic web were the ones like *Distributed Description Logics* [4] by Borgida et al., *E-connections* [25] by Kutz et al., *Context-OWL* [5] by Bouqet et al., and the work of CKR [30, 23] by Serafini et al. These were mainly logics based on DLs, which formalized contexts as OWL KBs, whose semantics is given using a distributed interpretation structure with additional semantic conditions that suits varying requirements. Compared to these works, the bridge rules we consider are much more expressive with conjunctions and existential variables that supports value/blank-node creation.

$\forall\exists$ *rules, TGDs, Datalog+- rules* Query answering over rules with universal existential quantifiers in the context of databases/KR, where these rules are called tuple generating dependencies (TGDs)/Datalog+- rules, was done by Beeri and Vardi [3] even in the early 80s, where the authors show that the query entailment problem in general is undecidable. However, recently many classes of such rules have been identified for which query answering is decidable. These includes (a) fragments s.t. the resulting models have bounded tree widths, called bounded treewidth sets (BTS), such as Weakly guarded rules [7], Frontier guarded rules [2], (b) fragments called finite unification sets (FUS), such as 'sticky' rules [6, 17], and (c) fragments called finite extension sets (FES), where sufficient conditions are enforced to ensure finiteness of the chase and its termination. The approach used for query answering in FUS is to rewrite the input query w.r.t. to the TGDs to another query that can be evaluated directly on the set of instances, s.t. the answers for the former query and latter query coincides. The approach is called the *query rewriting approach*. FES classes uses certain termination guarantying tests that check whether certain sufficient conditions are satisfied by the structure of TGDs. A large number of classes in FES are based on tests that detects 'acyclicity conditions' by analyzing the information flow between the TGD rules. *Weak acyclicity* [13, 11], was one of the first such notions, and was extended to *joint acyclicity* [24], *super weak acyclicity* [26], and *model faithful acyclicity* [9]. The most similar approach to ours is the weak acyclicity technique, where the structure of the rules is analyzed using a dependency graph that models the propagation of constants across various predicates positions, and restricting the dependency graph to be acyclic. Although this technique can be used in our scenario by translating a quad-system to a set of TGDs; if the obtained translation is weakly acyclic, then one could use existing algorithms for

Quad-System Fragment	dChase size w.r.t input quad-system	Data Complexity of CCQ entailment	Combined Complexity of CCQ entailment
Unrestricted Quad-Systems	Infinite	Undecidable	Undecidable
Context acylic Quad-Systems	Double exponential	PTIME-complete	2EXPTIME-complete

Table 1: Complexity info for various quad-system fragments

chase computation for the TGDs to compute the chase, the query entailment check can be done by querying the obtained chase. However, our approach has the advantage of straightforward implementability on existing quad-stores.

6 Summary and Conclusion

In this paper, we study the problem of query answering over contextualized RDF knowledge. We show that the problem in general is undecidable, and present a decidable class called context acyclic quad-systems. Table 1 summarizes the main results obtained. We can show that the notion of context acyclicity, introduced in section 4 can be used to extend the currently established tools for contextual reasoning to give support for expressive BRs with conjuction and existentials with decidability guarantees. We view the results obtained in this paper as a general foundation for contextual reasoning and query answering over contextualized RDF knowledge formats such as Quads, and can straightforwardly be used to extend existing Quad stores to encorporate for-all existential BRs of the form (1).

References

1. Abiteboul, S., Hull, R., Vianu, V.: Foundations of Databases. Addison-Wesley (1995)
2. Baget, J., Mugnier, M., Rudolph, S., Thomazo, M.: Walking the Complexity Lines for Generalized Guarded Existential Rules. In: IJCAI. pp. 712–717 (2011)
3. Beeri, C., Vardi, M.Y.: The Implication Problem for Data Dependencies. In: ICALP. pp. 73–85 (1981)
4. Borgida, A., Serafini, L.: Distributed Description Logics: Assimilating Information from Peer Sources. J. Data Semantics 1, 153–184 (2003)
5. Bouquet, P., Giunchiglia, F., van Harmelen, F., Serafini, L., Stuckenschmidt, H.: C-OWL: Contextualizing Ontologies. In: ISWC. pp. 164–179 (2003)
6. Calì, A., Gottlob, G., Pieris, A.: Query Answering under Non-guarded Rules in Datalog+/-. In: Hitzler, P., Lukasiewicz, T. (eds.) RR. Lecture Notes in Computer Science, vol. 6333, pp. 1–17. Springer (2010)
7. Cali, A., Gottlob, G., Lukasiewicz, T., Marnette, B., Pieris, A.: Datalog+/-: A Family of Logical Knowledge Representation and Query Languages for New Applications. In: Logic in Computer Science (LICS), 2010 25th Annual IEEE Symposium on. pp. 228 –242 (july 2010)
8. Carroll, J., Bizer, C., Hayes, P., Stickler, P.: Named graphs, provenance and trust. In: Proc. of the 14th int.l. conf. on WWW. pp. 613–622. ACM, New York, NY, USA (2005)

9. Cuenca Grau, B., Horrocks, I., Krötzsch, M., Kupke, C., Magka, D., Motik, B., Wang, Z.: Acyclicity Conditions and their Application to Query Answering in Description Logics. In: Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR'12). pp. 243–253. AAAI Press (2012)
10. Deutsch, A., Nash, A., Remmel, J.: The chase revisited. In: Proceedings of the twenty-seventh ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems. pp. 149–158. PODS '08 (2008)
11. Deutsch, A., Tannen, V.: Reformulation of XML Queries and Constraints. In: In ICDT. pp. 225–241 (2003)
12. D.Lenat: The Dimensions of Context Space. Tech. rep., CYCorp (1998), published online `http://www.cyc.com/doc/context-space.pdf`
13. Fagin, R., Kolaitis, P.G., Miller, R.J., Popa, L.: Data Exchange: Semantics and Query Answering. In: Theoretical Computer Science. pp. 28(1):89–124 (2005)
14. Ghidini, C., Serafini, L.: Distributed first order logics. In: Frontiers Of Combining Systems 2, Studies in Logic and Computation. pp. 121–140. Research Studies Press (1998)
15. Giunchiglia, F., Ghidini, C.: Local models semantics, or contextual reasoning = locality + compatibility. Artificial Intelligence 127 (2001)
16. Glimm, B., Lutz, C., Horrocks, I., Sattler, U.: Answering conjunctive queries in the $\mathcal{SHIQ}$ description logic. In: Proceedings of the IJCAI'07. pp. 299–404. AAAI Press (2007)
17. Gottlob, G., Manna, M., Pieris, A.: Polynomial Combined Rewritings for Existential Rules. In: KR'14: International Conference on Principles of Knowledge Representation and Reasoning (2014)
18. Harrison, M.A.: Introduction to Formal Language Theory. Addison-Wesley Longman Publishing Company, Inc., Boston, MA, USA, 1st edn. (1978)
19. Hayes, P. (ed.): RDF Semantics. W3C Recommendation (Feb 2004), `http://www.w3.org/TR/rdf-mt/`
20. ter Horst, H.J.: Completeness, decidability and complexity of entailment for RDF Schema and a semantic extension involving the OWL vocabulary. Web Semantics: Science, Services and Agents on the WWW 3(2-3), 79–115 (2005)
21. J.McCarthy: Generality in AI. Comm. of the ACM 30(12), 1029–1035 (1987)
22. Johnson, D.S., Klug, A.C.: Testing containment of conjunctive queries under functional and inclusion dependencies. Computer and System Sciences 28, 167–189 (1984)
23. Joseph, M., Serafini, L.: Simple reasoning for contextualized RDF knowledge. In: Proc. of Workshop on Modular Ontologies (WOMO-2011) (2011)
24. Krötzsch, M., Rudolph, S.: Extending decidable existential rules by joining acyclicity and guardedness. In: Walsh, T. (ed.) Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI'11). pp. 963–968. AAAI Press/IJCAI (2011)
25. Kutz, O., Lutz, C., Wolter, F., Zakharyaschev, M.: E-Connections of Abstract Description Systems. Artificial Intelligence 156(1), 1–73 (2004)
26. Marnette, B.: Generalized schema-mappings: from termination to tractability. In: Proceedings of the twenty-eighth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems. pp. 13–22. PODS '09, ACM, New York, NY, USA (2009)
27. McCarthy, J., Buvac, S., Costello, T., Fikes, R., Genesereth, M., Giunchiglia, F.: Formalizing Context (Expanded Notes) (1995)
28. R.Guha: Contexts: a Formalization and some Applications. Ph.D. thesis, Stanford (1992)
29. Schueler, B., Sizov, S., Staab, S., Tran, D.T.: Querying for meta knowledge. In: WWW '08: Proceeding of the 17th international conference on World Wide Web. pp. 625–634. ACM, New York, NY, USA
30. Serafini, L., Homola, M.: Contextualized knowledge repositories for the semantic web. Web Semantics: Science, Services and Agents on the World Wide Web (2012)

A Proofs of Section 3

Proof (Theorem 1). We show that CCQ entailment is undecidable for unrestricted quad-systems, by showing that the well known undecidable problem of "non-emptiness of intersection of context-free grammars" is reducible to the CCQ entailment problem.

Given an alphabet Σ, string $\boldsymbol{w}$ is a sequence of symbols from Σ. A language L is a subset of Σ^*, where Σ^* is the set of all strings that can be constructed from the alphabet Σ, and also includes the empty string ϵ. Grammars are machineries that generate a particular language. A grammar G is a quadruple $\langle V, T, S, P\rangle$, where V is the set of variables, T, the set of terminals, $S \in V$ is the start symbol, and P is a set of production rules (PR), in which each PR $r \in P$, is of the form:

$$\boldsymbol{w} \rightarrow \boldsymbol{w}'$$

where $\boldsymbol{w}, \boldsymbol{w}' \in \{T \cup V\}^*$. Intuitively application of a PR r of the form above on a string $\boldsymbol{w}_1$, replaces every occurrence of the sequence $\boldsymbol{w}$ in $\boldsymbol{w}_1$ with $\boldsymbol{w}'$. PRs are applied starting from the start symbol S until it results in a string $\boldsymbol{w}$, with $\boldsymbol{w} \in \Sigma^*$ or no more production rules can be applied on $\boldsymbol{w}$. In the former case, we say that $\boldsymbol{w} \in L(G)$, the language generated by grammar G. For a detailed review of grammars, we refer the reader to Harrison et al. [18]. A *context-free grammar* (CFG) is a grammar, whose set of PRs P, have the following property:

Property 1. For a CFG, every PR is of the form $v \rightarrow \boldsymbol{w}$, where $v \in V$, $\boldsymbol{w} \in \{T \cup V\}^*$.

Given two CFGs, $G_1 = \langle V_1, T, S_1, P_1\rangle$ and $G_2 = \langle V_2, T, S_2, P_2\rangle$, where V_1, V_2 are the set of variables, T s.t. $T \cap (V_1 \cup V_2) = \emptyset$ is the set of terminals. $S_1 \in V_1$ is the start symbol of G_1, and P_1 are the set of PRs of the form $v \rightarrow \boldsymbol{w}$, where $v \in V$, $\boldsymbol{w}$ is a sequence of the form $w_1 ... w_n$, where $w_i \in V_1 \cup T$. Similarly s_2, P_2 is defined. Deciding whether the language generated by the grammars $L(G_1)$ and $L(G_2)$ have non-empty intersection is known to be undecidable [18].

Given two CFGs, $G_1 = \langle V_1, T, S_1, P_1\rangle$ and $G_2 = \langle V_2, T, S_2, P_2\rangle$, we encode grammars G_1, G_2 into a quad-system of the form $QS_c = \langle Q_c, R\rangle$, with a single context identifier c. Each PR $r = v \rightarrow \boldsymbol{w} \in P_1 \cup P_2$, with $\boldsymbol{w} = w_1 w_2 w_3 .. w_n$, is encoded as a BR of the form:

$$c\colon (x_1, w_1, x_2), c\colon (x_2, w_2, x_3), ..., c\colon (x_n, w_n, x_{n+1}) \rightarrow c\colon (x_1, v, x_{n+1}) \qquad (7)$$

where $x_1, .., x_{n+1}$ are variables. W.l.o.g. we assume that the set of terminal symbols T is equal to the set of terminal symbols occurring in $P_1 \cup P_2$. For each terminal symbol $t_i \in T$, R contains a BR of the form:

$$c\colon (x, \texttt{rdf:type}, C) \rightarrow \exists y\, c\colon (x, t_i, y), c\colon (y, \texttt{rdf:type}, C) \qquad (8)$$

and Q_c contains only the triple:

$$c\colon (a, \texttt{rdf:type}, C)$$

We in the following show that:

$$QS_c \models \exists y\, c\colon (a, S_1, y) \wedge c\colon (a, S_2, y) \Leftrightarrow L(G_1) \cap L(G_2) \neq \emptyset$$

Claim. (1) For any $\boldsymbol{w} = t_1, ..., t_p \in T^*$, there exists $b_1, ...b_p$, s.t. $c\colon (a, t_1, b_1)$, $c\colon (b_1, t_2, b_2)$, ..., $c\colon (b_{p-1}, t_p, b_p)$, $c\colon (b_p, \texttt{rdf:type}, C) \in dChase(\ QS_c)$.

we proceed by induction on the $\|\boldsymbol{w}\|$.

base case suppose if $\|\boldsymbol{w}\| = 1$, then $\boldsymbol{w} = t_i$, for some $t_i \in T$. But Since by construction $c\colon (a, \texttt{rdf:type}, C) \in dChase_0(QS_c)$, on which rules of the form (8) is applicable. Hence, there exists an i s.t. $dChase_i(QS_c)$ contains $c\colon (a, t_i, b_i)$, $c\colon (b_i, \texttt{rdf:type}, C)$, for each $t_i \in T$. Hence, the base case.

hypothesis for any $\boldsymbol{w} = t_1...t_p$, if $\|\boldsymbol{w}\| \leq p'$, then there exists $b_1, ..., b_p$, s.t. $c\colon (a, t_1, b_1)$, $c\colon (b_1, t_2, b_2)$, ..., $c\colon (b_{p-1}, t_p, b_p)$, $c\colon (b_p,$ $\texttt{rdf:type}, C) \in dChase(QS_c)$.

inductive step suppose $\boldsymbol{w} = t_1...t_{p+1}$, with $\|\boldsymbol{w}\| \leq p' + 1$. Since $\boldsymbol{w}$ can be written as $\boldsymbol{w'}t_{p+1}$, where $\boldsymbol{w'} = t_1...t_p$, and by hypothesis, there exists $b_1, ..., b_p$ such that $c\colon (a, t_1, b_1)$, $c\colon (b_1, t_2, b_2)$, ..., $c\colon (b_{p-1}, t_p, b_p)$, $c\colon (b_p, \texttt{rdf:type}, C) \in dChase(QS_c)$. Also since rules of the form (8) are applicable on $c\colon (b_p, \texttt{rdf:type}, C)$, and hence produces triples of the form $c\colon (b_p, t_i, b^i_{p+1})$, $c\colon (b^i_{p+1}), \texttt{rdf:type}, C)$, for each $t_i \in T$. Since $t_{p+1} \in T$, the claim follows.

For a grammar $G = \langle V, T, S, P\rangle$, whose start symbol is S, and for any $\boldsymbol{w} \in \{V \cup T\}^*$, for some $V_j \in V$, we denote by $V_j \rightarrow^i \boldsymbol{w}$, the fact that $\boldsymbol{w}$ was derived from V_j by i production steps, i.e. there exists steps $V_j \rightarrow r_1, ..., r_i \rightarrow \boldsymbol{w}$, which lead to the production of $\boldsymbol{w}$. For any $\boldsymbol{w}$, $\boldsymbol{w} \in L(G)$, iff there exists an i s.t. $S \rightarrow^i \boldsymbol{w}$. For any $V_j \in V$, we use $V_j \rightarrow^* \boldsymbol{w}$ to denote the fact that there exists an arbitrary i, s.t. $V_j \rightarrow^i \boldsymbol{w}$.

Claim. (2) For any $\boldsymbol{w} = t_1...t_p \in \{V \cup T\}^*$, and for any $V_j \in V$, if $V_j \rightarrow^* \boldsymbol{w}$ and there exists $b_1, ..., b_{p+1}$, with $c\colon (b_1, t_1, b_2), ..., c\colon (b_p, t_p, b_{p+1}) \in dChase(QS_c)$, then $c\colon (b_1, V_j, b_{p+1}) \in dChase(QS_c)$.

We prove this by induction on the size of $\boldsymbol{w}$.

base case Suppose $\|\boldsymbol{w}\| = 1$, then $\boldsymbol{w} = t_k$, for some $t_k \in T$. If there exists b_1, b_2 s.t. $c\colon (b_1, t_k, b_2)$. But since there exists a PR $V_j \rightarrow t_k$, by transformation given in (7), there exists a BR $c\colon (x_1, t_k, x_2) \rightarrow c\colon (x_1, V_j, x_2) \in R$, which is applicable on $c\colon (b_1, t_k, b_2)$ and hence the quad $c\colon (b_1, V_j, b_2) \in dChase(QS_c)$.

hypothesis For any $\boldsymbol{w} = t_1...t_p$, with $\|\boldsymbol{w}\| \leq p'$, and for any $V_j \in V$, if $V_j \rightarrow^* \boldsymbol{w}$ and there exists $b_1, ...b_p, b_{p+1}$, s.t. $c\colon (b_1, t_1, b_2)$, ..., $c\colon (b_p, t_p, b_{p+1}) \in dChase(QS_c)$, then $c\colon (b_1, V_j, b_{p+1}) \in dChase(QS_c)$.

inductive step Suppose if $\boldsymbol{w} = t_1...t_{p+1}$, with $\|\boldsymbol{w}\| \leq p' + 1$, and $V_j \rightarrow^i \boldsymbol{w}$, and there exists $b_1, ...b_{p+1}, b_{p+2}$, s.t. $c\colon (b_1, t_1, b_2)$, ..., $c\colon (b_{p+1}, t_{p+1}, b_{p+2}) \in dChase(Q_c)$. Also, one of the following holds (i) $i = 1$, or (ii) $i > 1$. Suppose (i) is the case, then it is trivially the case that $c\colon (b_1, V_j, b_{p+2}) \in dChase(QS_c)$. Suppose if (ii) is the case, one of the two sub cases holds (a) $V_j \rightarrow^{i-1} V_k$, for some $V_k \in V$ and $V_k \rightarrow^1 \boldsymbol{w}$ or (b) there exist a $V_k \in V$, s.t. $V_k \rightarrow^* t_{q+1}...t_{q+l}$, with $2 \leq l \leq p$, where $V_j \rightarrow^* t_1...t_q V_k t_{p-l+1}...t_{p+1}$. If (a) is the case, trivially then $c\colon (b_1, V_k, b_{q+2}) \in dChase(QS_c)$, and since by construction there exists $c\colon (x_0, V_k, x_1) \rightarrow c\colon (x_0, V_{k+1}, x_1)$, ..., $c\colon (x_0, V_{k+i}, x_1) \rightarrow c\colon (x_0, V_j, x_1) \in R$, $c\colon (b_1, V_j, b_{q+2}) \in dChase(\ QS_c)$. If (b) is the case, then since $\|t_{q+1}...t_{q+l}\| \geq 2$, $\|t_1...t_q V_2 t_{p-l+1}...t_{p+1}\| \leq p'$. This implies that $c\colon (b_1, V_j, b_{p+2}) \in dChase(QS_c)$.

Similarly, by construction of $dChase(QS_c)$, the following claim can straightforwardly be shown to hold:

Claim. (3) For any $\boldsymbol{w} = t_1...t_p \in \{V \cup T\}^*$, and for any $V_j \in V$, if there exists $b_1, ..., b_p, b_{p+1}$, with $c\colon (b_1, t_1, b_2)$, ..., $c\colon (b_p, t_p, b_{p+1}) \in dChase(QS_c)$ and $c\colon (b_1, V_j, b_{p+1}) \in dChase(QS_c)$, then $V_j \rightarrow^* \boldsymbol{w}$.

(a) For any $\boldsymbol{w} = t_1...t_p \in T^*$, if $\boldsymbol{w} \in L(G_1) \cap L(G_2)$, then by claim 1, since there exists $b_1, ..., b_p$, s.t. $c\colon (a, t_1, b_1), ..., c\colon (b_{p-1}, t_p, b_p) \in dChase(QS_c)$. But since $\boldsymbol{w} \in L(G_1)$ and $\boldsymbol{w} \in L(G_2)$, $S_1 \rightarrow \boldsymbol{w}$ and $S_2 \rightarrow \boldsymbol{w}$. Hence by claim 2, $c\colon (a, S_1, b_p), c\colon (a, S_2, b_p) \in dChase(QS_c)$, which implies that $dChase(QS_c) \models \exists y\, c\colon (a, s_1, y) \wedge c\colon (a, s_2, y)$. Hence, $QS_c \models \exists y\, c\colon (a, s_1, y) \wedge c\colon (a, s_2, y)$.
(b) Suppose if $QS_c \models \exists y\ c\colon (a, S_1, y) \wedge c\colon (a, S_2, y)$, then this implies that there exists b_p s.t. $c\colon (a, S_1, b_p)$, $c\colon (a, S_2, b_p) \in dChase(QS_{\mathcal{C}})$. Then it is the case that there exists $\boldsymbol{w} = t_1...t_p \in T^*$, and $b_1, ..., b_p$ s.t. $c\colon (a, t_1, b_1)$, ..., $c\colon (b_{p-1}, t_p, b_p)$, $c\colon (a, S_1, b_p)$, $c\colon (a, S_2, b_p) \in dChase(QS_c)$. Then by claim 3, $S_1 \rightarrow^* \boldsymbol{w}$, $S_2 \rightarrow^* \boldsymbol{w}$. Hence, $w \in L(G_1) \cap L(G_2)$.

By (a),(b) it follows that there exists $\boldsymbol{w} \in L(G_1) \cap L(G_2)$ iff $QS_c \models \exists y\, c\colon (a, s_1, y) \wedge c\colon (a, s_2, y)$. As we have shown that the intersection of CFGs, which is an undecidable problem, is reducible to the problem of query entailment on unrestricted quad-system, the latter is undecidable.

B Proofs of Section 4

Proof (Theorem 2).

(i) (Membership) See lemma 4 for the membership in PTIME.

(Hardness) In order to prove P-hardness, we reduce a well known P-complete problem, 3HornSat, i.e. the satisfiability of propositional Horn formulas with at most 3 literals. Note that a (propositional) Horn formula is formula of the form:

$$P_1 \wedge \ldots \wedge P_n \rightarrow P_{n+1} \tag{9}$$

where P_i, for $1 \leq i \leq n+1$, are either propositional variables or constants t, f, that represents true and false, respectively. Note that for any propositional variable P, the fact that "P holds" is represented by the formula $t \rightarrow P$, and "P does not hold" is represented by the formula $P \rightarrow f$. A 3Horn formula is a formula of the form (9), where $1 \leq n \leq 2$. Note that any (set of) Horn formula(s) Φ can be transformed in polynomial time to a polynomially sized set Φ' of 3Horn formulas, by introducing auxiliary propositional variables s.t. Φ is satisfiable iff Φ' is satisfiable. A pure 3Horn formula is 3Horn formula of the form 9, where $n = 2$. Any 3Horn formula ϕ that is not pure can be trivially converted to equivalent pure form by appending a $\wedge\, t$ on the head part of ϕ. For instance, $P \rightarrow Q$, can be converted to $P \wedge t \rightarrow Q$. Hence, w.l.o.g. we assume that any set of 3Horn formulas is pure, and is of the form:

$$P_1 \wedge P_2 \rightarrow P_3 \tag{10}$$

We, in the following, reduce the satisfiability problem of pure 3Horn formulas to CCQ entailment problem over a quad-system whose set of schema triples, the set of BRs, and the CCQ CQ are all fixed.

For any set of pure Horn formulas Φ, we construct the quad-system $QS_{\mathcal{C}} = \langle Q_{\mathcal{C}}, R\rangle$, where $\mathcal{C} = \{c_t, c_f\}$. For any formula $\phi \in \Phi$ of the form (10), $Q_{\mathcal{C}}$ contains a quad $c_f\colon (P_1, P_2, P_3)$. In addition $Q_{\mathcal{C}}$ contains a quad $c_t\colon (t, \texttt{rdf:type}, T)$. R is the singleton that contains only the following fixed BR:

$$c_t\colon (x_1, \texttt{rdf:type}, T), c_t\colon (x_2, \texttt{rdf:type}, T), c_f\colon (x_1, x_2, x_3) \rightarrow c_t\colon (x_3, \texttt{rdf:type}, T)$$

Let the CQ be the fixed query $c_t(f, \texttt{rdf:type}, T)$.

Now, it is easy to see that $QS_{\mathcal{C}} \models CQ$, iff Φ is not satisfiable.

(ii) (Membership) See lemma 4.

(Hardness) See following heading. □

2EXPTIME-Hardness of CCQ Entailment In this subsection, we show that the combined complexity of the decision problem of CCQ entailment for context acyclic quad-systems is 2EXPTIME-hard. We show this by reduction of the word-problem of a 2EXPTIME deterministic turing machine (DTM) to the CCQ entailment problem. A DTM M is a tuple $M = \langle Q, \Sigma, \Delta, q_0, q_A\rangle$, where

- Q is a set of states,
- Σ is a finite alphabet that includes the blank symbol □,
- $\Delta\colon (Q \times \Sigma) \rightarrow (Q \times \Sigma \times \{+1, -1\})$ is the transition function,
- $q_0 \in Q$ is the initial state.
- $q_A \in Q$ is the accepting state.

W.l.o.g. we assume that there exists exactly one accepting state, which is also a halting state. A configuration is a word $\boldsymbol{\alpha} \in \Sigma^* Q \Sigma^*$. A configuration $\boldsymbol{\alpha}_2$ is a successor of the configuration $\boldsymbol{\alpha}_1$, iff one of the following holds:

1. $\boldsymbol{\alpha}_1 = \boldsymbol{w}_l q \sigma \sigma_r \boldsymbol{w}_r$ and $\boldsymbol{\alpha}_2 = \boldsymbol{w}_l \sigma' q' \sigma_r \boldsymbol{w}_r$, if $\Delta(q, \sigma) = (q', \sigma', R)$, or
2. $\boldsymbol{\alpha}_1 = \boldsymbol{w}_l q \sigma$ and $\boldsymbol{\alpha}_2 = \boldsymbol{w}_l \sigma' q' \square$, if $\Delta(q, \sigma) = (q', \sigma', R)$, or
3. $\boldsymbol{\alpha}_1 = \boldsymbol{w}_l \sigma_l q \sigma \boldsymbol{w}_r$ and $\boldsymbol{\alpha}_2 = \boldsymbol{w}_l q' \sigma_l \sigma' \boldsymbol{w}_r$, if $\Delta(q, \sigma) = (q', \sigma', L)$.

where $q, q' \in Q$, $\sigma, \sigma', \sigma_l, \sigma_r \in \Sigma$, and $\boldsymbol{w}_l, \boldsymbol{w}_r \in \Sigma^*$. Since number of configurations can at most be doubly exponential in the size of the input string, the number of tape cells traversed by the DTM tape head is also bounded double exponentially. A configuration $\boldsymbol{c} = \boldsymbol{w}_l q \boldsymbol{w}_r$ is an accepting configuration iff $q = q_A$. A language $L \subseteq \Sigma^*$ is accepted by a 2EXPTIME bounded DTM M, iff for every $\boldsymbol{w} \in L$, M accepts $\boldsymbol{w}$ in time $\mathcal{O}(2^{2^{\|\boldsymbol{w}\|}})$.

Simulating DTMs using Context Acyclic Quad-Systems Consider a DTM $M = \langle Q, \Sigma, \Delta, q_0, q_A\rangle$, and a string $\boldsymbol{w}$, with $\|\boldsymbol{w}\| = n$. Since the number of storage cells is double exponentially bounded, we first construct a quad-system $QS_{\mathcal{C}}^M = \langle Q_{\mathcal{C}}^M, R\rangle$,

where $\mathcal{C} = \{c_1, ..., c_n\}$, with $n = \|\boldsymbol{w}\|$. We follow the technique in works such as [6, 24] to iteratively generate a doubly exponential number of objects that represent the cells of the tape of the DTM. Let $Q_{\mathcal{C}}^{M}$ be initialized with the following quads:

$$c_0 \colon (k_0, \texttt{rdf:type}, R), c_0 \colon (k_1, \texttt{rdf:type}, R),$$
$$c_0 \colon (k_0, \texttt{rdf:type}, min_0), c_0 \colon (k_1, \texttt{rdf:type}, max_0), c_0 \colon (k_0, succ_0, k_1)$$

Now for each pair of elements of type R in c_i, a skolem blank-node is generated in c_{i+1}, and hence follows the recurrence relation $r(m+1) = [r(m)]^2$, with seed $r(1) = 2$, which after n iterations yields 2^{2^n}. In this way, a doubly exponential long chain of elements is created in c_n using the following set of rules:

$$c_i \colon (x_0, \texttt{rdf:type}, R), c_i \colon (x_1, \texttt{rdf:type}, R) \rightarrow$$
$$\exists y \; c_{i+1} \colon (x_0, x_1, y), c_{i+1} \colon (y, \texttt{rdf:type}, R)$$

The combination of minimal element with the minimal element (elements of type min_i) in c_i create the minimal element in c_{i+1}, and similarly the combination of maximal element with the maximal element (elements of type max_i) in c_i create the maximal element of c_{i+1}

$$c_{i+1} \colon (x_0, x_0, x_1), c_i \colon (x_0, \texttt{rdf:type}, min_i) \rightarrow c_{i+1} \colon (x_1, \texttt{rdf:type}, min_{i+1})$$
$$c_{i+1} \colon (x_0, x_0, x_1), c_i \colon (x_0, \texttt{rdf:type}, max_i) \rightarrow c_{i+1} \colon (x_1, \texttt{rdf:type}, max_{i+1})$$

Successor relation $succ_{i+1}$ is created in c_{i+1} using the following set of rules, using the well-known, integer counting technique:

$$c_i \colon (x_1, succ_i, x_2), c_{i+1} \colon (x_0, x_1, x_3), c_{i+1} \colon (x_0, x_2, x_4) \rightarrow c_{i+1} \colon (x_3, succ_{i+1}, x_4)$$

$$c_i \colon (x_1, succ_i, x_2), c_{i+1} \colon (x_1, x_3, x_5), c_{i+1} \colon (x_2, x_4, x_6), c_i \colon (x_3, \texttt{rdf:type}, max_i),$$
$$c_i \colon (x_4, \texttt{rdf:type}, min_i) \rightarrow c_{i+1} \colon (x_5, succ_{i+1}, x_6)$$

Each of the above set rules are instantiated for $0 \leq i < n$, and in this way after n generating dChase iterations, c_n has doubly exponential number of elements of type R, that are ordered linearly using the relation $succ_n$. By virtue of the first rule below, each of the objects representing the cells of the DTM are linearly ordered by the relation $succ$. Also the transitive closure of $succ$ is defined as the relation relation $succt$

$$c_n \colon (x_0, succ_n, x_1) \rightarrow c_n \colon (x_0, succ, x_1)$$
$$c_n \colon (x_0, succ, x_1) \rightarrow c_n \colon (x_0, succt, x_1)$$
$$c_n \colon (x_0, succt, x_1), c_n \colon (x_1, succt, x_2) \rightarrow c_n \colon (x_0, succt, x_2)$$

Also using a similar construction, we could create a linearly ordered chain of double exponential number of objects in c_n that represents configurations of M, whose minimal element is of type $conInit$, and the linear order relation being $conSucc$.

Various triple patterns that are used to encode the possible configurations, runs and their relations in M are:

$(x_0, head, x_1)$ denotes the fact that in configuration x_0, the head of the DTM is at cell x_1.

$(x_0, state, x_1)$ denotes the fact that in configuration x_0, the DTM is in state x_1.

(x_0, σ, x_1) where $\sigma \in \Sigma$, denotes the fact that in configuration x_0, the cell x_1 contains σ.

$(x_0, succ, x_1)$ denotes the linear order between cells of the tape.

$(x_0, succt, x_1)$ denotes the transitive closure of $succ$.

$(x_0, conSucc, x_1)$ to denote the fact that x_1 is a successor configuration of x_0.

$(x_0, \texttt{rdf:type}, Accept)$ denotes the fact that the configuration x_0 is an accepting configuration.

Since in our construction, each $\sigma \in \Sigma$ is represented as relation, we could constrain that no two alphabets $\sigma \neq \sigma'$ are on the same cell using the following axiom:

$$c_n\colon (z_1, \sigma, z_2), c_n\colon (z_1, \sigma', z_2) \rightarrow$$

for each $\sigma \neq \sigma' \in \Sigma$. Note that the above BR has an empty head, is equivalent to asserting the negation of its body.

Initialization Suppose the initial configuration is $q_0\boldsymbol{w}\Box$, where $\boldsymbol{w} = \sigma_0 ... \sigma_{n-1}$, then we enforce this using the following BRs in our quad-system $QS^M_{\mathcal{C}}$ as:

$$c_n\colon (x_0, \texttt{rdf:type}, conInit), c_n\colon (x_1, \texttt{rdf:type}, min_n) \rightarrow c_n\colon (x_0, head, x_1),$$
$$c_n\colon (x_0, state, q_0)$$
$$c_n\colon (x_0, \texttt{rdf:type}, min_n) \wedge \bigwedge_{i=0}^{n-1} c_n\colon (x_i, succ, x_{i+1}) \wedge c_n\colon (x_j, \texttt{rdf:type},$$
$$conInit) \rightarrow \bigwedge_{i=0}^{n-1} c_n\colon (x_j, \sigma_i, x_i) \wedge c_n\colon (x_j, \Box, x_n)$$
$$c_n\colon (x_j, \texttt{rdf:type}, conInit), c_n\colon (x_j, \Box, x_0), c_n\colon (x_0, succt, x_1) \rightarrow c_n\colon (x_j, \Box, x_1)$$

The last BR copies the $\Box$ to every succeeding cell in the initial configuration.

Transitions For every left transition $\Delta(q, \sigma) = (q_j, \sigma', -1)$, the following BR:

$$c_n\colon (x_0, head, x_i), c_n\colon (x_0, \sigma, x_i), c_n\colon (x_0, state, q), c_n\colon (x_j, succ, x_i), c_n\colon (x_0,$$
$$conSucc, x_1) \rightarrow c_n\colon (x_1, head, x_j), c_n\colon (x_1, \sigma', x_i), c_n\colon (x_1, state, q_j)$$

For every right transition $\Delta(q, \sigma) = (q_j, \sigma', +1)$, the following BR:

$$c_n\colon (x_0, head, x_i), c_n\colon (x_0, \sigma, x_i), c_n\colon (x_0, state, q), c_n\colon (x_i, succ, x_j), c_n\colon (x_0,$$
$$conSucc, x_1), \rightarrow c_n\colon (x_1, head, x_j), c_n\colon (x_1, \sigma', x_i), c_n\colon (x_1, state, q_j)$$

Inertia If in any configuration the head is at cell i of the tape, then in every successor configuration, elements in preceding and following cells of i in the tape are retained. The following two BRs ensures this:

$$c_n\colon (x_0, head, x_i), c_n\colon (x_0, conSucc, x_1), c_n\colon (x_j, succt, x_i), c_n\colon (x_0, \sigma, x_j) \\ \rightarrow c_n\colon (x_1, \sigma, x_j)$$

$$c_n\colon (x_0, head, x_i), c_n\colon (x_0, conSucc, x_1), c_n\colon (x_i, succt, x_j), c_n\colon (x_0, \sigma, x_j) \\ \rightarrow c_n\colon (x_1, \sigma, x_j)$$

The rules above are instantiated for every $\sigma \in \Sigma$.

Acceptance An configuration whose state is q_A is accepting:

$$c_n\colon (x_0, state, q_A) \rightarrow c_n\colon (x_0, \texttt{rdf:type}, Accept)$$

If a configuration of accepting type is reached, then it can be back propagated to the initial configuration, using the following BR:

$$c_n\colon (x_0, conSucc, x_1), c_n\colon (x_1, \texttt{rdf:type}, Accept) \rightarrow c_n\colon (x_0, \texttt{rdf:type}, Accept)$$

Finally since M accepts $\boldsymbol{w}$ iff the initial configuration is an accepting configuration, i.e. $QS_{\mathcal{C}}^{M} \models \exists y\ c_n\colon (y, \texttt{rdf:type}, conInit), c_n\colon (y, \texttt{rdf:type}, Accept)$. Since there is no edge from any c_j to c_i, for each $1 \leq i < j \leq n$, the context dependency graph for $QS_{\mathcal{C}}^{M}$ is acyclic, and hence $QS_{\mathcal{C}}^{M}$ is context acyclic. Since we reduced the word problem of 2EXPTIME DTM, which is a 2EXPTIME-hard problem, to CCQ entailment problem over context acyclic quad-systems, it immediately follows that CCQ entailment problem over context acyclic quad-systems is 2EXPTIME-hard. □